

## 2 | A CRASH IN THE DESERT

On a desert road near Florence, Arizona, one morning in the fall of 2005, a Volkswagen Touareg was kicking up a dust cloud, bouncing along at a steady twenty to twenty-five miles per hour, carrying four passengers. To the casual observer there was nothing unusual about the way the approaching vehicle was being driven. The road was particularly rough, undulating up and down through a landscape dotted with cactus and scrubby desert vegetation. The car bounced and wove, and all four occupants were wearing distinctive crash helmets. The Touareg was plastered with decals like a contestant in the Baja 1000 off-road race. It was also festooned with five curious sensors perched at the front of the roof, each with an unobstructed view of the road. Other sensors, including several radars, also sprouted from the roof. A video camera peered out through the windshield. A tall whip antenna set at the back of the vehicle, in combination with the sensors, conspired to give a postapocalyptic vibe reminiscent of a Mad Max movie.

The five sensors on the roof were actually mechanical contraptions, each rapidly sweeping an infrared laser beam back and forth over the road ahead. The beams, invisible to the eye, constantly reflected off the gravel road and the desert surrounding the vehicle. Bouncing back to sensors, the lasers provided a constantly changing portrait of the surrounding landscape accurate to the centimeter. Even small rocks on the road hundreds of feet ahead could not escape the unblinking gaze of the sensors, known as lidar.

The Touareg was even more peculiar inside. The driver, Sebastian Thrun, a roboticist and artificial intelligence researcher, wasn't driving. Instead he was gesturing with his hands as he chatted with the other passengers. His eyes rarely watched the road. Most striking of all: his hands never touched the steering wheel, which twitched back and forth as if controlled by some unseen ghost.

Sitting behind Thrun was another computer researcher, Mike Montemerlo, who wasn't driving either. His eyes were buried in the screen of a laptop computer that was displaying the data from the lasers, radars, and cameras in a God's-eye view of the world around the car in which potential obstacles appeared as a partial rainbow of blips on a radar screen. It revealed an ever-changing cloud of colored dots that in aggregate represented the road unfolding ahead in the desert.

The car, named Stanley, was being piloted by an ensemble of software programs running on five computers installed in the trunk. Thrun was a pioneer of an advanced version of a robotic navigation technique known as SLAM, which stands for simultaneous localization and mapping. It had become a standard tool for robots to find their way through previously unexplored terrain. The wheel continued to twitch back and forth as the car rolled along the rutted road lined with cactus and frequent outcroppings of boulders. Immediately to Thrun's right, between the front seats, was a large red E-Stop button to override the car's autopilot in an emergency. After a half-dozen miles, the robotic meanderings of the Touareg felt anticlimactic. Stanley wasn't driving down the freeway, so as the desert scenery slid by, it seemed increasingly unnecessary to wear crash

helmets for what was more or less a Sunday drive in the country.

The car was in training to compete in the Pentagon's second Grand Challenge, an ambitious autonomous vehicle contest intended to jump-start technology planned for future robotic military vehicles. At the beginning of the twenty-first century, Congress instructed the U.S. military to begin designing autonomous vehicles. Congress even gave the Pentagon a specific goal: by 2015, one-third of the army's vehicles were supposed to go places without human drivers present. The directive wasn't clear as to whether both autonomous and remotely teleoperated vehicles would satisfy the requirement. In either case the idea was that smart vehicles would save both money and soldiers' lives. But by 2004, little progress had been made, and Tony Tether, the then controversial director of the Pentagon's blue-sky research arm, DARPA, the Defense Advanced Research Projects Agency, came up with a high-profile contest as a gambit to persuade computer hackers, college professors, and publicity-seeking corporations to innovate where the military had failed. Tether was a product of the military-industrial complex, and the contest itself was a daring admission that the defense contracting world was not able to get the job done. By opening the door for ragtag teams of hobbyists, Tether ran the risk of undermining the classified world dominated by the Beltway Bandits that surround Washington, D.C., and garner the lion's share of military research dollars.

The first Grand Challenge contest, held in 2004, was something of a fiasco. Vehicles tipped over, drove in circles, and ignominiously knocked down fences. Even the most successful entrant had gotten stuck in the dust just seven miles from the starting line in a 120-mile race, with one wheel spinning helplessly as it teetered off the edge of the road. When the dust settled, a reporter flying overhead in a light plane saw brightly colored vehicles scattered motionless over the desert floor. At the time it seemed obvious that self-driving cars were still years away, and Tether was criticized for organizing a publicity stunt.

Now, just a little more than a year later, Thrun was behind the wheel in a second-generation robot contestant. It felt like the future had arrived sooner than expected. It took only a dozen miles, however, to realize that techno-enthusiasm is frequently premature. Stanley crested a rise in the desert and plunged smartly into a swale. Then, as the car tilted upward, its laser guidance system swept across an overhanging tree limb. Without warning the robot navigator spooked, the car wrenched violently first left, then right, and instantly plunged off the road. It all happened faster than Thrun could reach over and pound the large red E-Stop button.

Luckily, the car found a relatively soft landing. The Touareg had been caught by an immense desert thornbush just off the road. It cushioned the crash landing and the car stopped slowly enough that the air bags didn't deploy. When the occupants surveyed the road from the crash scene, it was obvious that it could have been much worse. Two imposing piles of boulders bracketed the bush, but the VW had missed them.

The passengers stumbled out and Thrun scrambled up on top of the vehicle to reposition the sensors bent out of alignment by the crash. Then everyone piled back into Stanley, and Montemerlo removed an offending block of software code that had been intended to make the ride more comfortable for human passengers. Thrun restarted the autopilot and the machine once again headed out into the Arizona desert. There were other mishaps that day, too. The AI controller had no notion of the consequence of mud puddles and later in the day Stanley found itself ensnared in a small lake in the middle of the road. Fortunately there were several human-driven support

vehicles nearby, and when the car's wheels began spinning helplessly, the support team of human helpers piled out to push the car out of the goo.

These were small setbacks for Thrun's team, a group of Stanford University professors, VW engineers, and student hackers among more than a dozen teams competing for a multimillion-dollar cash prize. The day was a low point after which things improved dramatically. Indeed, the DARPA contest would later prove to be a dividing line between a world in which robots were viewed as toys or research curiosities and one in which people began to accept that robots could move about freely.

Stanley's test drive was a harbinger of technology to come. The arrival of machine intelligence had been forecast for decades in the writings of science-fiction writers, so much so that when the technology actually began to appear, it seemed anticlimactic. In the late 1980s, anyone wandering through the cavernous Grand Central Station in Manhattan would have noticed that almost a third of the morning commuters were wearing Sony Walkman headsets. Today, of course, the Walkmans have been replaced by Apple's iconic bright white iPhone headphones, and there are some who believe that technology haute couture will inevitably lead to a future version of Google Glass—the search engine maker's first effort to augment reality—or perhaps more ambitious and immersive systems. Like the frog in the pot, we have been desensitized to the changes wrought by the rapid increase and proliferation of information technology.

The Walkman, the iPhone, and Google Glass all prefigure a world where the line between what is human and who is machine begins to blur. William Gibson's *Neuromancer*, the science-fiction novel that popularized the idea of cyberspace, drew a portrait of a new cybernetic territory composed of computers and networks. It also painted a future in which computers were not discrete boxes, but would be woven together into a dense fabric that was increasingly wrapped around human beings, "augmenting" their senses.

It is not such a big leap to move from the early-morning commuters wearing Sony Walkman headsets, past the iPhone users wrapped in their personal sound bubbles, directly to Google Glass—wearing urban hipsters watching tiny displays that annotate the world around them. They aren't yet "jacked into the net," as Gibson foresaw, but it is easy to assume that computing and communication technology is moving rapidly in that direction.

Gibson was early to offer a science-fiction vision of what has been called "intelligence augmentation." He imagined computerized inserts he called "microsofts"—with a lowercase *m*—that could be snapped into the base of the human skull to instantly add a particular skill—like a new language. At the time—several decades ago—it was obviously an impossible bit of science fiction. Today his cyborg vision is something less of a wild leap.

In 2013 President Obama unveiled the BRAIN initiative, an effort to simultaneously record the activities of one million neurons in the human brain. But one of the major funders of the BRAIN initiative is DARPA, and the agency is not interested in just reading *from* the brain. BRAIN scientists will patiently explain that one of the goals of the plan is to build a *two-way* interface between the human brain and computers. On its face, such an idea seems impossibly sinister, conjuring up images of the ultimate Big Brother and thought control. At the same time there is a utopian implication inherent in the technology. The potential future is perhaps the inevitable trajectory of human-computer interaction design, implicit in J. C. R. Licklider's 1960 manifesto,

“Man-Computer Symbiosis,” where he foretold a more intimate collaboration between humans and machines.

While the world of *Neuromancer* was wonderful science fiction, actually entering the world that Gibson portrayed presents a puzzle. On one hand, the arrival of cyborgs poses the question of what it means to be human. By itself that isn't a new challenge. While technology may be evolving increasingly rapidly today, humans have always been transformed by technology, as far back as the domestication of fire or the invention of the wheel (or its eventual application to luggage in the twentieth century). Since the beginning of the industrial era machines have displaced human labor. Now with the arrival of computing and computer networks, for the first time machines are displacing “intellectual” labor. The invention of the computer generated an earlier debate over the consequences of intelligent machines. The new wave of artificial intelligence technologies has now revived that debate with a vengeance.

Mainstream economists have maintained that over time the size of the workforce has continued to grow despite the changing nature of work driven by technology and innovation. In the nineteenth century, more than half of all workers were engaged in agricultural labor; today that number has fallen to around 2 percent—and yet there are more people working than ever in occupations outside of agriculture. Indeed, even with two recessions, between 1990 and 2010 the overall workforce in the United States increased by 21 percent. If the mainstream economists are correct, there is no economic cataclysm on a societal level due to automation in the offing.

However, today we are entering an era where humans can, with growing ease, be designed in or out of “the loop,” even in formerly high-status, high-income, white-collar professional areas. On one end of the spectrum smart robots can load and unload trucks. On the other end, software “robots” are replacing call center workers and office clerks, as well as transforming high-skill, high-status professions such as radiology. In the future, how will the line be drawn between man and machine, and who will draw it?

Despite the growing debate over the consequences of the next generation of automation, there has been very little discussion about the designers and their values. When pressed, the computer scientists, roboticists, and technologists offer conflicting views. Some want to replace humans with machines; some are resigned to the inevitability—“I for one, welcome our insect overlords” (later “robot overlords”) was a meme that was popularized by *The Simpsons*—and some of them just as passionately want to build machines to extend the reach of humans. The question of whether true artificial intelligence—the concept known as “Strong AI” or Artificial General Intelligence—will emerge, and whether machines can do more than mimic humans, has also been debated for decades. Today there is a growing chorus of scientists and technologists raising new alarms about the possibility of the emergence of self-aware machines and their consequences. Discussions about the state of AI technology today veer into the realm of science fiction or perhaps religion. However, the reality of machine autonomy is no longer merely a philosophical or hypothetical question. We have reached the point where machines are capable of performing many human tasks that require intelligence as well as muscle: they can do factory work, drive vehicles, diagnose illnesses, and understand documents, and they can certainly control weapons and kill with deadly accuracy.

The AI versus IA dichotomy is nowhere clearer than in a new generation of weapons systems now on the horizon. Developers at DARPA are about to cross a new technological threshold with a

replacement for today's cruise missiles, the Long Range Anti-Ship Missile, or LRASM. Developed for the navy, it is scheduled for the U.S. fleet in 2018. Unlike its predecessors, this is a new weapon in the U.S. arsenal with the ability to make targeting decisions autonomously. The LRASM is designed to fly to an enemy fleet while out of contact with human controllers and then use artificial intelligence technologies to decide which target to kill.

The new ethical dilemma is, will humans allow their weapons to pull triggers on their own without human oversight? Variations of that same challenge are inherent in rapid computerization of the automobile, and indeed transportation in general is emblematic of the consequences of the new wave of smart machines. Artificial intelligence is poised to have an impact on society that will be greater than the effect that personal computing and the Internet have had beginning in the 1990s. Significantly, the transformation is being shepherded by a group of elite technologists.

Several years ago Jerry Kaplan, a Silicon Valley veteran who began his career as a Stanford artificial intelligence researcher and then became one of those who walked away from the field during the 1980s, warned a group of Stanford computer scientists and graduate student researchers: "Your actions today, right here in the Artificial Intelligence Lab, as embodied in the systems you create, may determine how society deals with this issue." The imminent arrival of the next generation of AI is a crucial ethical challenge, he contended: "We're in danger of incubating robotic life at the expense of our own life."<sup>1</sup> The dichotomy that he sketched out for the researchers is the gap between intelligent machines that displace humans and human-centered computing systems that extend human capabilities.

Like many technologists in Silicon Valley, Kaplan believes we are on the brink of the creation of an entire economy that runs largely without human intervention. That may sound apocalyptic, but the future Kaplan described will almost certainly arrive. His deeper point was that today's technology acceleration isn't arriving blindly. The engineers who are designing our future are each—individually—making choices.

On an abandoned military base in the California desert during the fall of 2007 a short, heavysset man holding a checkered flag stepped out onto a dusty makeshift racing track and waved it energetically as a Chevrolet Tahoe SUV glided past at a leisurely pace. The flag waver was Tony Tether, the director of DARPA.

There was no driver behind the wheel of the vehicle, which sported a large GM decal. Closer examination revealed no passengers in the car, and none of the other cars in the "race" had drivers or passengers either. Viewing the event, in which the cars glided seemingly endlessly through a makeshift town previously used for training military troops in urban combat, it didn't seem to be a race at all. It felt more like an afternoon of stop-and-go Sunday traffic in a science-fiction movie like *Blade Runner*.

Indeed, by almost any standard it was an odd event. The DARPA Urban Challenge pitted teams of roboticists, artificial intelligence researchers, students, automotive engineers, and software hackers against each other in an effort to design and build robot vehicles capable of driving autonomously in an urban traffic setting. The event was the third in the series of contests that Tether organized. At the time military technology largely amplified a soldier's killing power rather than replacing the soldier. Robotic military planes were flown by humans and, in some

cases, by extraordinarily large groups of soldiers. A report by the Defense Science Board in 2012 noted that for many military operations it might take a team of several hundred personnel to fly a single drone mission.<sup>2</sup>

Unmanned ground vehicles were a more complicated challenge. The problem in the case of ground vehicles was, as one DARPA manager would put it, that “the ground was hard”—“hard” as in “hard to drive on,” rather than as in “rock.” Following a road is challenging enough, but robot car designers are confronted with an endless array of special cases: driving at night, driving into the sun, driving in rain, on ice—the list goes on indefinitely.

Consider the problem of designing a machine that knows how to react to something as simple as a plastic bag in a lane on the highway. Is the bag hard, or is it soft? Will it damage the vehicle? In a war zone, it might be an improvised explosive device. Humans can see and react to such challenges seemingly without effort, when driving at low speed with good visibility. For AI researchers, however, solving that problem is the holy grail in computer vision. It became one of a myriad of similar challenges that DARPA set out to solve in creating the autonomous vehicle Grand Challenge events. In the 1980s roboticists in both Germany and the United States had made scattered progress toward autonomous driving, but the reality was that it was easier to build a robot to go to the moon than to build one that could drive by itself in rush-hour traffic. And so Tony Tether took up the challenge. The endeavor was risky: if the contests failed to produce results, the series of Grand Challenge self-driving contests would become known as Tether’s Folly. Thus the checkered flag at the final race proved to be as much a victory lap for Tether as for the cars.

There had been darker times. Under Tether’s directorship the agency hired Admiral John Poindexter to build the system known as Total Information Awareness. A vast data-mining project that was intended to hunt terrorists online by collecting and connecting the dots in oceans of credit card, email, and phone records, the project started a privacy firestorm and was soon canceled by Congress in May of 2003. Although Total Information Awareness vanished from public view, it in fact moved into the nation’s intelligence bureaucracy only to become visible again in 2013 when Edward Snowden leaked hundreds of thousands of documents that revealed a deep and broad range of systems for surveillance of any possible activity that could be of interest. In the pantheon of DARPA directors, Tether was also something of an odd duck. He survived the Total Information Awareness scandal and pushed the agency ahead in other areas with a deep and controlling involvement in all of the agency’s research projects. (Indeed, the decision by Tether to wave the checkered flag was emblematic of his tenure at DARPA—Tony Tether was a micromanager.)

DARPA was founded in response to the Soviet Sputnik, which was like a thunderbolt to an America that believed in its technological supremacy. With the explicit mission of ensuring the United States was never again technologically superseded by another power, the directors of DARPA—at birth more simply named the Advanced Research Projects Agency—had been scientists and engineers willing to place huge bets on blue-sky technologies, with close relationships and a real sense of affection for the nation’s best university researchers.

Not so with Tony Tether, who represented the George W. Bush era. He had worked for decades as a program manager for secretive military contractors and, like many surrounding George W. Bush, was wary of the nation’s academic institutions, which he thought were too independent to be trusted with the new mission. Small wonder. Tether’s worldview had been formed when he was an

electrical engineering grad student at Stanford University during the 1960s, where there was a sharp division between the antiwar students and the scientists and engineers helping the Vietnam War effort by designing advanced weapons.

After arriving as director he went to work changing the culture of the agency that had gained a legendary reputation for the way it helped invent everything from the Internet to stealth fighter technology. He rapidly moved money away from the universities and toward classified work done by military contractors supporting the twin wars in Iraq and Afghanistan. The agency moved away from “blue sky” toward “deliverables.” Publicly Tether made the case that it was still possible to innovate in secret, as long as you fostered the competitive culture of Silicon Valley, with its turmoil of new ideas and rewards for good tries even if they failed.

And Tether certainly took DARPA in new technology directions. His concern for the thousands of maimed veterans coming back without limbs and with increasing the power and effectiveness of military decision-makers inspired him to push agency dollars into human augmentation projects as well as artificial intelligence. That meant robotic arms and legs for wounded soldiers, and an “admiral’s advisor,” a military version of what Doug Engelbart had set out to do in the 1960s with his vision of intelligence augmentation, or IA. The project was referred to as PAL, for Perceptive Assistant that Learns, and much of the research would be done at SRI International, which dubbed the project CALO, or Cognitive Assistant that Learns and Organizes.

It was ironic that Tether returned to the research agenda originally promoted during the mid-1960s by two visionary DARPA program managers, Robert Taylor and J. C. R. Licklider. It was also bittersweet, although few mentioned it, that despite Doug Engelbart’s tremendous early success in the early 1970s, his project had faltered and fallen out of favor at SRI. He ended up being shuffled off to a time-sharing company for commercialization, where his project sat relatively unnoticed and underfunded for more than a decade. The renewed DARPA investment would touch off a wave of commercial innovation—CALO would lead most significantly to Apple’s Siri personal assistant, a direct descendant of the augmentation approach originally pioneered by Engelbart.

Tether’s automotive Grand Challenge drew garage innovators and eager volunteers out of the woodwork. In military terms it was a “force multiplier,” allowing the agency to get many times the innovation it would get from traditional contracting efforts. At its heart, however, the specific challenge that Tether chose to pursue had been cooked up more than a decade earlier inside the same university research community that he now disfavored. The guiding force behind the GM robot SUV that would win the Urban Challenge in 2007 was a Carnegie Mellon roboticist who had been itching to win this prize for more than a decade.

In the fall of 2005, Tether’s second robot race through the California desert had just ended at the Nevada border and Stanford University’s roboticists were celebrating. Stanley, the once crash-prone computerized Volkswagen Touareg, had just pulled off a come-from-behind victory and rolled under a large banner before a cheering audience of several thousand.

Just a few feet away in another tent, however, the atmosphere had the grim quality of a losing football team’s locker room. The Carnegie Mellon team had been the odds-on favorite, with two robot vehicle entries and a no-nonsense leader, a former marine and rock climber, William L. “Red” Whittaker. His team had lost the race due to a damnable spell of bad luck. Whittaker had



barnstormed into the first DARPA race eighteen months earlier with another heavily funded GM Humvee, only to fail when the car placed a wheel just slightly off road on a steep climb. Trapped in the sand, it was out of the competition. Up to then, Whittaker's robot had been head and shoulders above the others. So when he returned the second time with a two-car fleet and a squad of photo analysts to pore over the course ahead of the competition, he had easily been cast as the odds-on favorite.

Once again, however, bad luck struck. His primary vehicle led Stanley until late in the race, when it malfunctioned, slowing dramatically and allowing the Stanford team to sail by and grab the \$2 million prize. After the second loss, Whittaker stood in the tent in front of his team and gave an inspiring speech worthy of any college football coach. "On any Sunday . . ." he told his team, echoing the words of the losing coach. The loss was especially painful because the leaders of the Stanford team, Sebastian Thrun and Mike Montemerlo, were former CMU roboticists who had defected to Stanford, where they organized the rival, winning effort. Years later the loss still rankled. Outside of Whittaker's office at the university is a portrait of the ill-fated team of robot car designers. In the hallway Whittaker would greet visitors and replay the failure in detail.

The defeat was particularly striking because Red Whittaker had in many ways been viewed widely as the nation's premier roboticist. By the time of the Grand Challenges he had already become a legend for designing robots capable of going places where humans couldn't go. For decades he combined a can-do attitude with an adventurer's spirit. His parents had both flown planes with a bit of barnstorming style. His father, an air force bomber pilot, sold mining explosives after the war. His mother, a chemist, was a pilot, too. When he was a young man, she had once flown him under a bridge.<sup>3</sup>

His Pennsylvania upbringing led him to develop a style of robotics that pushed in the direction of using the machines primarily as tools to extend an adventurer's reach, a style in the tradition of Yvon Chouinard, the legendary climber who designed and made his own climbing hardware, or Jacques Cousteau, the undersea explorer who made his own breathing equipment. With a degree in civil engineering from Princeton and a two-year tour as a marine sergeant, the six-foot-four Whittaker pioneered "field" robotics—building machines that left the laboratory and moved around in the world.

In Red Whittaker's robotic world, however, humans were still very much in the loop. In every case he used them to extend his reach as an adventurer. He had built machines used in nuclear power plant catastrophes at both Three Mile Island and Chernobyl. In the late 1980s he designed a huge nineteen-foot-tall robot called Ambler that was intended to walk on Mars. He sent a robot into a volcano and had been one of the first roboticists in the United States to explore the idea of an autonomous car as part of Carnegie Mellon's Navlab project.

"This is not the factory of the future," he was fond of pointing out. "The ideas that make it in the factory don't make it in the outside world."<sup>4</sup>

As a young man Whittaker had variously been a rower, wrestler, boxer, and mountain climber. His love of adventure had not been without personal pain, however. He spent a decade of his life rock climbing, sneaking away from his robot projects to spend time in Yosemite and the Himalayas. He even soloed the east wall of the Matterhorn in winter conditions. He had begun climbing casually as a member of a local explorer's club in Pittsburgh. It only became a passion when he met another young climber, after seeing a notice on a bulletin board: "Expert climber



willing to teach the right guy,” the note read, adding: “You must have a car.”

The two would become inseparable climbing partners over the next decade.

That magic time for Whittaker came to a sudden end one summer when they were climbing in Peru. His Pittsburgh friend was climbing with another young climber. The two were roped together and the younger man slipped and pulled both men down a tumbling set of ledges for almost a thousand feet. Whittaker, who was off-rope during the accident, was able to rescue the young climber, but his friend was killed by the fall. Whittaker returned to Pittsburgh shaken by the accident. It would take months before he mustered up the courage to go over to the home where the young man had lived with his parents and clean out the dead climber’s room.

The death left its mark. Whittaker stopped climbing, but still hungered for some sort of challenging adventure. He began to build ever more exotic robots, capable of performing tasks ranging from simple exploration to sophisticated repair, to extend his adventures into volcanoes, and ultimately, perhaps, to the moon and Mars. Even when he had been climbing on Earth in the 1970s and 1980s, it was becoming more and more difficult to find virgin territory. With the possibility of “virtual exploration,” new vistas would open up indefinitely and Whittaker could again dream of climbing and rappelling, this time perhaps with a humanoid robot stand-in on another world.

Whittaker redeemed his bitter loss to Stanford’s Stanley several years later in the third Grand Challenge, in 2007. His General Motors–backed “Boss” would win the final Urban Driving Challenge.

One of the most enduring bits of Silicon Valley lore recalls how Steve Jobs recruited Pepsi CEO John Sculley to Apple by asking him if he wanted to spend the rest of his life selling sugar water. Though some might consider it naive, the Valley’s ethos is about changing the world. That is at the heart of the concept of “scale,” which is very much a common denominator in motivating the region’s programmers, hardware hackers, and venture capitalists. It is not enough to make a profit, or to create something that is beautiful. It has to have an impact. It has to be something that goes under 95 percent of the world’s Christmas trees, or offers clean water or electricity to billions of people.

Google’s chief executive Larry Page took the Steve Jobs approach in recruiting Sebastian Thrun. Thrun was a fast-rising academic who had spent a sabbatical year at Stanford in 2001, which opened his eyes to the world that Silicon Valley offered beyond the walls of academia. There was more out there besides achieving tenure, publishing, and teaching students.

He returned to Stanford as an assistant professor in 2003. He attended the first DARPA Grand Challenge as an observer. The self-driving car competition completely changed his perspective: he realized that there were great thinkers outside of his cloistered academic community who cared deeply about changing the world. In between, during his short return to CMU, he had sent a note to Whittaker offering to help their software effort, but was rebuffed. Thrun had brought a group of students with him from CMU when he returned to Stanford, including Mike Montemerlo, whose father was a NASA roboticist. Montemerlo gave a presentation on the first DARPA contest. At the end of his presentation his final slide asked, “Should we at Stanford enter the Grand Challenge?” And then he answered his own question in a large font. “NO!” There were a dozen reasons not to do it. They would have no chance of winning, it was too hard, it would cost too much money.

Thrun looked at Montemerlo and it was obvious that although on paper he was the quintessential pessimist, everything in his demeanor was saying yes.



Sebastian Thrun (*left*) and Mike Montemerlo (*right*) in front of the Stanford University autonomous vehicle while it was being tested to take part in DARPA's Urban Challenge in 2007. (*Photo courtesy of the author*)

Soon afterward Thrun threw himself into the DARPA competition with passion. For the first time in his life he felt like he was focusing on something that was genuinely likely to have broad impact. Living in the Arizona desert for weeks on end, surviving on pizza, the team worked on the car until it was able to drive the backcountry roads flawlessly.

Montemerlo and Thrun made a perfect team of opposites. Montemerlo was fundamentally conservative, and Thrun was extraordinarily risk-inclined. As head of software, Montemerlo would build his conservative assumptions into his programs. When he wasn't looking, Thrun would go through the code and comment out the limitations to make the car go faster. It would infuriate the younger researcher. But in the end it was a winning combination.

Larry Page had said to Thrun that if you really focus on something you can achieve amazing things. He was right. After Stanley captured the \$2 million DARPA prize, Thrun took Page's words to heart. The two men had become friends after Thrun helped the Google cofounder debug a home robot that Page had been tinkering with. Thrun borrowed the device and brought it back able to navigate inside Page's home.

Navigation, a necessity for autonomous robots, had become Thrun's particular expertise. At CMU and later at Stanford he worked to develop SLAM, the mapping technique pioneered at Stanford Research Institute by the designers of the first mobile robots beginning in the 1960s. Thrun had helped make the technique fast and accurate and had paved the way for using it in autonomous cars. At Carnegie Mellon he had begun to attract national attention for a variety of mobile robots. In 1998 at the Smithsonian in D.C., he showcased Minerva, a mobile museum tour guide that was connected to the Web and could interact with museum guests and travel up to three and a half miles per hour. He worked with Red Whittaker to send robots into mines, which relied heavily on SLAM techniques. Thrun also tried to integrate mobile and autonomous robots in

nursing and elder-care settings, with little success. It turned out to be a humbling experience, which gave him a deep appreciation of the limitations of using technologies to solve human problems. In 2002, in a team effort between the two universities, Thrun pioneered a new flavor of SLAM that was dubbed FastSLAM, which could be used in real-world situations where it was necessary to locate thousands of objects. It was an early example of a new wave of artificial intelligence and robotics that increasingly relied on probabilistic statistical techniques rather than on rule-based inference.

At Stanford, Thrun would rise quickly to become director of the revitalized Stanford Artificial Intelligence Laboratory that had originally been created by John McCarthy in the 1960s. But he also quickly became frustrated by the fragmented life of an academic, dividing time between teaching, public speaking, grant writing, working on committees, doing research, and mentoring. In the wake of his 2005 DARPA Grand Challenge victory Thrun had also become more visible in high-technology circles. His talks described the mass atrocities committed by human drivers that resulted in more than one million killed and maimed each year globally. He personalized the story. A close friend had been killed in an automobile accident when Thrun was a high school student in his native Germany. Many people he was close to lost friends in accidents. More recently, a family member of a Stanford faculty secretary was crippled for life after a truck hit her car. In an instant she went from being a young girl full of life and possibility to someone whose life was forever impaired. Thrun's change-the-world goals gave him a platform at places like the TED Conference.

After building two vehicles for the DARPA Challenge contests, he decided to leave Stanford. Page offered him the opportunity to do things at "Google scale," which meant that his work would touch the entire world. He secretly set up a laboratory modeled vaguely on Xerox PARC, the legendary computer science laboratory that was the birthplace of the modern personal computer, early computer networks, and the laser printer, creating projects in autonomous cars and reinventing mobile computing. Among other projects, he helped launch Google Glass, which was an effort to build computing capabilities including vision and speech into ordinary glasses.

Unlike laboratories of the previous era that emphasized basic science, such as IBM Research and Bell Labs, Google's X Lab was closer in style to PARC, which had been established to vault the copier giant, restyled "the Document Company," into the computer industry—to compete directly with IBM. The X Lab was intended to push Google into new markets. Google felt secure in its Web search monopoly so, with a profit stream that by the end of 2013 was more than \$1 billion a month, the search company funded ambitious R & D projects that might have nothing to do with the company's core business. Google was famous for its 70-20-10 rule, which gave its engineers free time to pursue their own side projects. Employees are supposed to spend 10 percent of their time on projects entirely unrelated to the company's core business. Its founders Sergey Brin and Larry Page believed deeply in thinking big. They called their efforts "moon shots": not pure science, but research projects that were hopefully destined to have commercial rather than purely scientific impact.

It was a perfect environment for Thrun. His first project in 2008 had been to create the company's fleet of Street View cars that systematically captured digital images of homes and businesses on every street in the nation. The next year he began an even more ambitious effort: a self-driving car that would travel on public streets and highways. He was both cautious and bold in

the car project. A single accident might destroy the Google car, so at the outset he ensured that a detailed safety regime was in place. He was acutely aware that if there was any indication in the program that Google had not been incredibly careful, it would be a disaster. He never let an untrained driver near the wheel of the small Toyota Prius fleet on which the system was being developed. The cars would eventually drive more than a half-million miles without an accident, but Thrun understood that even a single error every fifty thousand to a hundred thousand miles was too high an error rate. At the same time he believed that there was a path forward that would allow Google to redefine what it meant to be in a car.

Like the automotive industry, Thrun and his team believed in the price/volume curve, which suggested that costs would go down the more a company manufactured a particular thing. Sure, today a single experimental lidar laser radar might cost tens of thousands of dollars, but the Google engineers had faith that in a few years it would be so cheap that it would not be a showstopper in the bill of materials of some future car. In the trade-off between cost and durability, Thrun always felt it would make sense to design and build more reliable systems now and depend on mass manufacturing technologies for price reductions to kick in later. The pricey laser guidance systems didn't actually contain that many parts, so there was little reason to believe that prices couldn't come down rapidly. It had already happened with radar, which had once been an esoteric military and aviation technology but in recent years had begun showing up in motion detectors and premium automobiles.

Thrun evinced an engineer's worldview and tended toward a libertarian outlook. He held a pro-business point of view that the global corporation was an evolutionary step beyond the nation-state. He also subscribed to the belief, commonplace in the Valley, that within three decades as much as 90 percent of all jobs will be made obsolete by advancing AI and robotic technologies. Indeed, Thrun believed that most people's jobs are actually pretty useless and unfulfilling. There are countless manual labor jobs—everything from loading and unloading trucks to driving them—that could vanish over the coming decade. He also believed that much of the bureaucratic labor force is actively counterproductive. Those people make other people's work harder. Thrun had a similar contempt for what he perceived as Detroit's hidebound car industry that could have easily used technology to radically reshape transportation systems and make them safer, but did little and was content to focus on changing the shape of a car's tail fins each year. By 2010 he had a deep surprise in store for an industry that did not change easily and was largely unfamiliar with Silicon Valley culture.<sup>5</sup>

The DARPA races created ripples in Detroit, the cradle of the American automotive industry, but the industry kept to its traditional position that cars were meant to be driven by people and should not operate autonomously. By and large the industry had generally resisted computer technology. Many car manufacturers adhered to a "computers are buggy" philosophy. However, engineers elsewhere in the country were beginning to think about transportation through the new lens of cheap sensors, the microprocessor, and the Internet.

In the spring of 2010, rumors about an experimental Google car began to float around Silicon Valley. Initially they sounded preposterous. The company, nominally a provider of Internet search, was supposedly hiding the cars in plain sight. Google engineers, so the story went, had

succeeded in robotically driving from San Francisco to Los Angeles on freeways at night! The notion immediately elicited both guffaws and pointed reminders that such an invention would be illegal, even if it was possible. How could they get away with something so crazy?

Of course, Google's young cofounders Sergey Brin and Larry Page had by then perfected a public image for wild schemes based on AI and other futuristic technologies to transform the world. Eric Schmidt, the company's chief executive officer beginning in 2001, would tell reporters that his role was one of adult supervision—persuading the cofounders which of their ideas should be kept above and which below the “bar.” The cofounders famously considered the idea of a space elevator. New, incredibly strong material had recently been developed, and this material was so strong that, rather than using a rocket, it would be possible to build a cable that reached from the Earth into orbit to inexpensively hoist people and materials into space. When queried about the idea Schmidt would pointedly state that this was one of the ideas that was being considered, but was—for the moment at least—“below the bar.”

In the hothouse community of technical workers that is Silicon Valley, however, it is difficult to keep secrets. It was obvious that something was afoot. Within a year after the final DARPA Grand Challenge event in 2007, Sebastian Thrun had taken a leave from Stanford and gone to work full-time at Google. His departure was never publicly announced, or even mentioned in the press, but among the Valley's digerati, Thrun's change of venue was noted with intense interest. A year later, while he was with colleagues in a bar in Alaska at an artificial intelligence conference, he spilled out a few tantalizing comments. Those words circulated back in Silicon Valley and made people wonder.

In the end, however, it was a high school friend of one of the low-paid drivers the company had hired to babysit its robotic Prius fleet who inadvertently spilled the beans. *One of the kids I went to high school with is being paid fifteen dollars an hour by Google to sit in a car while it drives itself!* a young college student blurted to me. At that point the secret became impossible to contain. The company was parking its self-driving cars in the open lots on the Google campus.

The Google engineers had made no effort to conceal the sensors attached to the roof of the ungainly-looking creatures, which looked even odder than their predecessor, Stanford's Stanley. Rather than an array of sensors mounted above the windshield, each Prius had a single 360-degree lidar, mounted a foot above the center of the car's roof. The coffee-can-sized mechanical laser, made by Velodyne, a local high-tech company, made it possible to easily create a real-time map of the surrounding environment for several hundred feet in all directions. It wasn't cheap—at the time the lidar alone added \$70,000 to the vehicle cost.

How did the odd-looking Toyotas, also equipped with less obtrusive radars, cameras, GPS, and inertial guidance sensors, escape discovery for as long as they did? There were several reasons. The cars were frequently driven at night, and the people who saw them confused them with a ubiquitous fleet of Google Street View cars, which had a large camera on a mast above the roof taking photographs that were used to build a visual map of the surrounding street as the car drove. (They also recorded people's Wi-Fi network locations, which then could be used as beacons to improve the precision in locating Google's Android smartphones.)

The Street View assumption usually hid the cars in plain sight, but not always. The Google engineer who had the pleasure of the first encounter with law enforcement was James Kuffner, a former CMU roboticist who had been one of the first members of the team. Kuffner had made a

name for himself at Carnegie Mellon working on both navigation and a variety of humanoid robot projects. His expertise was in motion planning, figuring out how to teach machines to navigate in the real world. He was bitten by the robot car bug as part of Red Whittaker's DARPA Grand Challenge team, and when key members of that group began to disappear into a secret Google project code-named Chauffeur, he jumped at the chance.

Late one night they were testing the robotic Prius in Carmel, one of the not-quite-urban driving areas they were focusing on closely. They were testing the system late at night because they were anxious to build detailed maps with centimeter accuracy, and it was easier to get baseline maps of the streets when no one was around. After passing through town several times with their distinctive lidar prominently displayed, Kuffner was sitting in the driver's seat when the Prius was pulled over by a local policeman suspicious about the robot's repeated passes.

"What is this?" he asked, pointing to the roof.

Kuffner, like all of the Google drivers, had been given strict instructions how to respond to this inevitable confrontation. He reached behind him and handed a prewritten document to the officer. The police officer's eyes widened as he read it. Then he grew increasingly excited and kept the Google engineers chatting late into the night about the future of transportation.

The incident did not lead to public disclosure, but once I discovered the cars in the company's parking lots while reporting for the *New York Times*, the Google car engineers relented and offered me a ride.

From a backseat vantage point it was immediately clear that in the space of just three years, Google had made a significant leap past the cars of the Grand Challenge. The Google Prius replicated much of the original DARPA technology, but with more polish. Engaging the autopilot made a whooshing *Star Trek* sound. Technically, the ride was a remarkable tour de force. A test drive began with the car casually gliding away from Google's campus on Mountain View city streets. Within a few blocks, the car had stopped at both stop signs and stoplights and then merged onto rush-hour traffic on the 101 freeway. At the next exit the car then drove itself off the freeway onto a flyover overpass that curved gracefully over the 101. What was most striking to the first-time passenger was the car's ability to steer around the curve exactly as a human being might. There was absolutely nothing robotic about AI's driving behavior.

When the *New York Times* published the story, the Google car struck Detroit like a thunderbolt. The automobile industry had been adding computer technology and sensors to cars at a maddeningly slow pace. Even though cruise control had been standard for decades, intelligent cruise control—using sensors to keep pace with traffic automatically—was still basically an exotic feature in 2010. A number of automobile manufacturers had outposts in Silicon Valley, but in the wake of the publicity surrounding the Google car, the remaining carmakers rushed to build labs close by. Nobody wanted to see a repeat of what happened to personal computer hardware makers when Microsoft Windows became an industry standard and hardware manufacturers found that their products were increasingly low-margin commodities while much of the profit in the industry flowed to Microsoft. The automotive industry now realized that it was facing the same threat.

At the same time, the popular reaction to the Google car was mixed. There had long been a rich science-fiction tradition of *Jetsons*-like futuristic robot cars. They had even been the stuff of TV series like *Knight Rider*, a 1980s show featuring a crime fighter assisted by an artificially



intelligent car. There was also a dark-side vision of automated driving, perhaps best expressed in Daniel Suarez's 2009 sci-fi thriller *Daemon*, in which AI-controlled cars not only drove themselves, but ran people down as well. Still, the general perception was a deep well of skepticism about whether driverless cars would ever become a reality. However, Sebastian Thrun had made his point abundantly clear that humans are terrible drivers, largely the consequence of human fallibility and inattention. By the time his project was discovered, Google cars had driven more than a hundred thousand miles without an accident, and over the next several years that number would rise above a half-million miles. A young Google engineer, Anthony Levandowski, routinely commuted from Berkeley to Mountain View, a distance of fifty miles, in one of the Priuses, and Thrun himself would let a Google car drive him from Mountain View to his vacation home in Lake Tahoe on weekends.

Today, partially autonomous cars are already appearing on the market, and they offer two paths toward the future of transportation—one with smarter and safer human drivers and one in which humans will become passengers.

Google had not disclosed how it planned to commercialize its research, but by the end of 2013 more than a half-dozen automakers had already publicly stated their intent to offer autonomous vehicles. Indeed, 2014 was the year that the line was first crossed commercially when a handful of European car manufacturers including BMW, Mercedes, Volvo, and Audi announced an optional feature—traffic jam assist, the first baby step toward autonomous driving. In Audi's case, while on the highway, the car will drive autonomously when traffic is moving at less than forty miles per hour, staying in its lane and requiring driver intervention only as dictated by lawyers fearful that passengers might go to sleep or otherwise distract themselves. In late 2014 Tesla announced that it would begin to offer an "autopilot" system for its Model S, making the car self-driving in some highway situations.

The autonomous car will sharpen the dilemma raised by the AI versus IA dichotomy. While there is a growing debate over the liability issue—who will pay when the first human is killed by a robot car—the bar that the cars must pass to improve safety is actually incredibly low. In 2012 a National Highway Transportation Safety Administration study estimated that the deployment of electronic stability control (ESC) systems in light vehicles alone would save almost ten thousand lives and prevent almost a quarter million injuries.<sup>6</sup> Driving, it would seem, might be one area of life where humans should be taken out of the loop to the greatest degree possible. Even unimpaired humans are not particularly good drivers, and we are worse when distracted by the gadgets that increasingly surround us. We will be saved from ourselves by a generation of cheap cameras, radars, and lidars that, when coupled with pattern-sensing computers, will wrap an all-seeing eye around our cars, whether we are driving or are being driven.

For Amnon Shashua, the aha moment came while seated in a university library as a young computer science undergraduate in Jerusalem. Reading an article written in Hebrew by Shimon Ullman, who had been the first Ph.D. student under David Marr, a pioneer in vision research, he was thrilled to discover that the human retina was in many ways a computer. Ullman was a computer scientist who specialized in studying vision in both humans and machines. The

realization that computing was going on inside the eye fascinated Shashua and he decided to follow in Ullman's footsteps.

He arrived at MIT in 1996 to study artificial intelligence when the field was still recovering from an earlier cycle of boom-and-bust. Companies had tried to build commercial expert systems based on the rules and logic approach of early artificial intelligence pioneers like Ed Feigenbaum and John McCarthy. In the heady early days of AI it had seemed that it would be straightforward to simply bottle the knowledge of a human expert, but the programs were fragile and failed in the marketplace, leading to the collapse of a number of ambitious start-ups. Now the AI world was rebounding. Progress in AI, which had been relatively stagnant for its first three decades, finally took off during the 1990s because statistical techniques made classification and decision-making practical. AI experiments hadn't yet seen great results because the computers of the era were still relatively underpowered for the data at hand. The new ideas, however, were in the air.

As a graduate student Shashua would focus on a promising approach to visually recognizing objects based on imaging them from multiple views to capture their geometry. The approach was derived from the world of computer graphics, where Martin Newell had pioneered a new modeling approach as a graduate student at the University of Utah—which was where much of computer graphics was invented during the 1970s. A real Melitta teapot found in his kitchen inspired Newell's approach. One day, as he was discussing the challenges of modeling objects with his wife over tea, she suggested that he model that teapot, which thereafter became an iconic image in the early days of computer graphics research.

At MIT, Shashua studied under computer vision scientists Tommy Poggio and Eric Grimson. Poggio was a scientist who stood between the worlds of computing and neuroscience and Grimson was a computer scientist who would later become MIT's chancellor. At the time there seemed to be a straight path from capturing shapes to recognizing them, but programming the recognition software would actually prove daunting. Even today the holy grail of "scene understanding"—for example, not only identifying a figure as a woman but also identifying what she might be doing—is still largely beyond reach, and significant progress has been made only in niche industries. For example, many cars can now identify pedestrians or bicyclists in time to automatically slow before a collision.

Shashua would become one of the masters in pragmatically carving out those niches. In an academic world where brain scientists debated computational scientists, he would ally himself with a group who took the position that "just because airplanes don't flap their wings, it doesn't mean they can't fly." After graduate school he moved back to Israel. He had already founded a successful company, Cognitens, using vision modeling to create incredibly accurate three-dimensional models of parts for industrial applications. The images, accurate to hair-thin tolerances, gave manufacturers ranging from automotive to aerospace the ability to create digital models of existing parts, enabling checking their fit and finish. The company was quickly sold.

Looking around for another project, Shashua heard from a former automotive industry customer about an automaker searching for stereovision technology for computer-assisted driving. They knew about Shashua's work in multiple-view geometry and asked if he had ideas for stereovision. He responded, "Well, that's fine but you don't need a stereo system, you can do it with a single camera." Humans can tell distances with one eye shut under some circumstances, he pointed out.

The entrepreneurial Shashua persuaded General Motors to invest \$200,000 to develop demonstration software. He immediately called a businessman friend, Ziv Aviram, and proposed that they start a new company. “There is an opportunity,” he told his friend. “This is going to be a huge field and everybody is thinking about it in the wrong way and we already have a customer, somebody who is willing to pay money.” They called the new company Mobileye and Shashua wrote software for the demonstration on a desktop computer, soon showing one-camera machine vision that seemed like science fiction to the automakers at that time.

Six months after starting the project, Shashua heard from a large auto industry supplier that General Motors was about to offer a competitive bid for a way to warn drivers that the vehicle was straying out of its lane. Until then Mobileye had been focusing on far-out problems like vehicle and pedestrian detection that the industry thought weren’t solvable. However, the parts supplier advised Shashua, “You should do something now. It’s important to get some real estate inside the vehicle, then you can build more later.”

The strategy made sense to Shashua, and so he put one of his Hebrew University students on the project for a couple of months. The lane-keeping software demonstration wasn’t awful, but he realized it probably wasn’t as good as what companies who’d started earlier could show, so there was virtually no way that the fledgling company would win.

Then he had a bright idea. He added vehicle detection to the software, but he told GM that the capability was a bug and that they shouldn’t pay attention. “It will be taken out in the next version, so ignore it,” he said. That was enough. GM was ecstatic about the safety advance that would be made possible by the ability to detect other vehicles at low cost. The automaker immediately canceled the bidding and committed to fund the novice firm’s project developments. Vehicle detection would facilitate a new generation of safety features that didn’t replace drivers, but rather augmented them with an invisible sensor and computer safety net. Technologies like lane departure warning, adaptive cruise control, forward collision warning, and anticollision braking are now rapidly moving toward becoming standard safety systems on cars.

Mobileye would grow into one of the largest international suppliers of AI vision technology for the automotive industry, but Shashua had bigger ideas. After creating Cognitens and Mobileye, he took a postdoctoral year at Stanford in 2001 and shared an office with Sebastian Thrun. Both men would eventually pioneer autonomous driving. Shashua would pursue the same technologies as Thrun, but with a more pragmatic, less “moon shot” approach. He had been deeply influenced by Poggio, who pursued biological approaches to vision, which were alternatives to using the brute force of increasingly powerful computers to recognize objects.

The statistical approach to computing would ultimately work best when both powerful clusters of computers, such as Google’s cloud, and big data sets were available. But what if you didn’t have those resources? This is where Shashua would excel. Mobileye had grown to become a uniquely Israeli technology firm, located in Jerusalem, close to Hebrew University, where Shashua teaches computer science. A Mobileye-equipped Audi served as a rolling research platform. Unlike the Google car, festooned with sensors, from the outside the Mobileye Audi looked normal, apart from a single video camera mounted unobtrusively just in front of the rearview mirror in the center of the windshield. The task at hand—automatic driving—required powerful computers, hidden in the car’s trunk, with some room left over for luggage.

Like Google, Mobileye has significant ambitions that are still only partially realized. On a

spring afternoon in 2013, two Mobileye engineers, Gaby Hayon and Eyal Bagon, drove me several miles east of Jerusalem on Highway 1 until they pulled off at a nondescript turnout where another employee waited in a shiny white Audi A7. As we got in the A7 and prepared for a test drive, Gaby and Eyal apologized to me. The car was a work in progress, they explained. Today Mobileye supplies computer vision technology to automakers like BMW, Volvo, Ford, and GM for safety applications. The company's third-generation technology is touted as being able to detect pedestrians and cyclists. Recently, Nissan gave a hint of things to come, demonstrating a car that automatically swerved to avoid a pedestrian walking out from behind a parked car.

Like Google, the Israelis are intent on going further, developing the technology necessary for autonomous driving. But while Google might decide to compete with the automobile industry by partnering with an upstart like Tesla, Shashua is exquisitely sensitive to the industry culture exemplified by its current customers. That means that his vision system designs must cost no more than several hundred dollars for even a premium vehicle and less than a hundred for a standard Chevy.

Google and Mobileye have taken very different approaches to solving the problem of making a car aware of its surroundings with better-than-human precision at highway speeds. Google's system is based on creating a remarkably detailed map of the world around the car using radars, video, and a Velodyne lidar, all at centimeter accuracy, augmenting the data it collects using its Street View vehicles. The Google car connects to the map database via a wireless connection to the Google cloud. The network is an electronic crutch for the car's navigation system, confirming what the local sensors are seeing around the car.

The global map database could make things easier for Google. One of the company's engineers confided that when the project got under way the Google team was surprised to find how dynamic the world is. Not only do freeway lanes frequently come and go for maintenance reasons, but "whole bridges will move," he said. Even without the database, the Google car is able to do things that might seem to be the province of humans alone, such as seamlessly merging into highway traffic and handling stop-and-go traffic in a dense urban downtown.

Google has conducted its project with a mix of Thrun's German precision and the firm's penchant for secrecy. The Israelis are more informal. On that hot spring afternoon in suburban Jerusalem there was little caution on the part of the Mobileye engineers. "Why don't you drive?" Eyal suggested to me, as he slid into the passenger seat behind a large display and keyboard. The engineers proceeded to give a rapid-fire minute-long lesson on driving a robot car: You simply turn on cruise control and then add the lane-keeping feature—steering—by pulling the cruise control stick on the steering wheel toward you. A heads-up display projected on the windshield showed the driver the car's speed and an icon indicated that the autonomous driving feature was on.

Unlike the Google car, which originally had a distinctive *Star Trek* start-up sound, there is only a small visual cue when the autopilot is engaged and the Mobileye Audi takes off down the highway by itself, at times reaching speeds of more than sixty miles per hour. On the road that snakes down a desolate canyon to the Dead Sea, it is difficult to relax. In an automated car, it is very challenging for a novice driver when, before long, a car ahead begins to slow for a stoplight. It takes all of one's willpower to keep a foot off the brake and trust the car as, sure enough, it slows down and rolls smoothly to a stop behind the vehicle ahead.

The Google car conveys the detached sense of a remote and slightly spooky machine intelligence at work somewhere in the background or perhaps in some distant cloud of computers. By contrast, during its test phase in 2013, the Mobileye car left a passenger acutely aware of the presence of machine assistance. The car needs to weave a bit within the lane when it starts to pull away from a stop—not a behavior that inspires confidence. If you understand the underlying technology, however, it ceases to be alarming. The Audi’s vision system uses a single “monocular” camera. The third dimension, depth, is computed based on a clever algorithm that Shashua and his researchers designed, referred to as “structure from motion,” and by weaving slightly the car is able to build a 3-D map of the world ahead of it.

Knowing that did little to comfort a first-time passenger, however. During the test ride, as it passed a parked car, the Audi pulled in the direction of the vehicle. Not wanting to see what the car was “really thinking,” I grabbed the wheel and nudged the Audi back into the center of the lane. The Israeli engineers showed no signs of alarm, only amusement. After a half hour’s drive along an old road that felt like it was still a part of antiquity, the trip was over. The autonomous ride felt jarringly like science fiction, yet it was just the first hint of what will gradually become a broad societal phase change. “Traffic jam assist” has already appeared on the market. Technology that was remarkable to someone visiting Israel in 2013 routinely handles stop-and-go freeway traffic around the world today.

The next phase of automatic driving will start arriving well before 2020—vehicles will handle routine highway driving, not just in traffic jams, but also in the commute from on-ramp to off-ramp. General Motors now calls this capability “Super Cruise,” and it will mark a major step in the changing role of humans as drivers—from manual control to supervision.

The Google vision is clearly to build a vehicle in which the human becomes a passenger and no longer participates in driving. Yet Shashua believes that even for Google the completely driverless vehicle is still far in the future. Such a car will stumble across what he describes as the four-way stop puzzle—a completely human predicament. At intersections without stoplights there is an elaborate social dance that goes on between drivers, and it will be difficult for independent, noncommunicating computer systems to solve anytime in the foreseeable future.

Another complication is that human drivers bend the rules and ignore the protocols frequently, and pedestrians add massive complications. These challenges may be the real barrier for a future of completely AI-based automobiles in urban settings: we haven’t yet been able to wrap our heads around the legal trouble posed by a possible AI-caused accident. There is a middle ground, Shashua believes, somewhere short of the Google vision, but realistic enough that it will begin to take over highway driving in just a couple of years. His approach wraps increasingly sophisticated sensor arrays and AI software around the human driver, who stays in the loop, a human with superawareness who can see farther and more clearly, and perhaps switch back and forth from other tasks besides driving. The car can alert the driver when his or her participation is beneficial or necessary, depending on the driver’s preferences. Or perhaps the car’s preferences.

Standing beside the Audi in the Jerusalem suburbs, it was clear that this was the new Promised Land. Like it or not, we are no longer in a biblical world, and the future is not about geographical territory, but rather about a rapidly approaching technological wonder world. Machines that will begin as golems are becoming brilliant and capable of performing many human tasks, from physical labor to rocket science.

Google had a problem. More than three years into the company's driverless car program, the small research team at the Mountain View-based Internet search company had safely driven more than a half-million miles autonomously. They had made startling progress in areas that had been generally believed impossible within the traditional automotive industry. Google cars could drive during the day and at night, could change lanes, and could even navigate the twisty Lombard Street in San Francisco. Google managed these advances by using the Internet to create a virtual infrastructure. Rather than building "smart" highways that would entail vast costs, they used the precise maps of the world created by the Google Street View car fleet.

Some of the achievements demonstrated an eerie humanlike quality. For example, the car's vision system had the ability to recognize construction zones, slow down accordingly, and make its way through the cones safely. It also could adjust for vehicles partially blocking a lane, moving over as necessary. The system had not only been able to recognize bicyclists, but it could identify their hand signals and slow down to allow for them to change lanes in front. That suggested that Google was making progress on an even harder problem: What would a driverless car do when it was confronted by a cop making hand signal motions at an accident or a construction zone?

MIT roboticist John Leonard had taken particular joy in driving around Cambridge and shooting videos of the most confounding situations for autonomous vehicles. In one of his videos his car rolls up to a stop sign at a T-intersection and is waiting to make a left turn. The car is delayed by a long line of traffic passing from right to left, without a stop sign. The situation is complicated by light traffic coming from the opposite direction. The challenge is to persuade the drivers in the slow lane to give way, while not colliding with one of the cars zipping by at higher speed in the other direction.<sup>7</sup>

The video that was perhaps the toughest challenge for the Google vision system was taken at a busy crosswalk somewhere downtown. There is a crowd of people at a pedestrian crossing with a stoplight. The car is approaching when suddenly, completely ignoring the green light for the cars, a police officer's hand abruptly shoots out on the left of the frame to stop traffic for the pedestrians. Ultimately it may not be an impossible problem to solve with computer vision. If today's systems can already recognize cyclists and hand signals, uniforms cannot be far behind. But it will not be solved easily or quickly.

Intent on transforming education with massive open online courses, or MOOCs, and not wishing to compete for leadership of X Lab with Google cofounder Sergey Brin, Thrun largely departed the research program in 2012. As is often the case in Silicon Valley, Thrun had not been able to see his project through. After creating and overseeing the secret X Laboratory at Google for several years, he decided it was time for him to move on when Brin joined the effort. Brin proposed that they be codirectors, but Thrun realized that with Google's cofounder in the mix he would no longer be in control and so it was time for a new challenge.

In the fall of 2011 Thrun and Peter Norvig had taught one of several free Stanford online course offerings, an Introduction to Artificial Intelligence. It made a big splash. More than 160,000 students signed up for the course, which was almost ten times the size of Stanford's nonvirtual student body. Although only a fraction of those enrolled in the course would ultimately complete it, it became a global "Internet moment": Thrun and Norvig's class raised the specter of a new low-cost form of education that would not only level the playing field by putting the



world's best teachers within reach of anyone in the world, but also threaten the business models of high-priced elite universities. Why pay Stanford tuition if you could take the course anyway as a City College student?

Thrun was still nominally participating one day a week at Google, but the project leadership role was taken by Chris Urmson, a soft-spoken roboticist who had been Red Whittaker's chief lieutenant in the DARPA vehicle challenges. He had been one of the first people that Thrun hired after he came to Google to start the then secret car program. In the summer of 2014 he said he wanted to create a reliable driverless car before his son reached driving age, which was about six years in the future.

After Thrun departed, Urmson took the program a long way toward its original goal of autonomous driving on the open road. Google had divided the world into highway driving and driving in urban conditions. At a press conference called to summarize their achievements, Google acknowledged that their greatest challenge was figuring out how to program the car to drive in urban areas. Urmson, however, argued in a posting he made on the company's website that the chaos of the city streets with cars, bicyclists, and pedestrians moving in apparently random fashion, was actually reasonably predictable. The Google training experiment had encountered thousands of these situations and the company had developed software models that would expect both the expected (a car stopped at a red light) and the unexpected (a car running a red light). He and his team implied that the highway driving challenge was largely solved, with one caveat—the challenge of keeping the human driver engaged. That problem presented itself when the Google team farmed out some of their fleet of robotic vehicles to Google employees to test during their daily commute. “We saw some things that made us nervous,” Urmson told a reporter. The original Google driving program had involved two professional drivers who worked from an aircraft-style checklist. The person in the driver's seat was vigilant and ready to take command if anything anomalous should happen. The real world was different. Some of the Google employees, on their way home after a long day at the office, had the disturbing habit of becoming distracted, up to and including falling asleep!

This was called the “handoff” problem. The challenge was to find a way to quickly bring a distracted human driver who might be reading email, watching a movie, or even sleeping back to the level of “situational awareness” necessary in an emergency. Naturally, people nodded off a lot more often in driverless cars that they had come to trust. It was something that the automotive industry would address in 2014 in the traffic jam assist systems that would drive cars in stop-and-go highway traffic. The drivers had to keep at least one hand on the wheel, except for ten-second intervals. If the driver didn't demonstrate “being there,” the system gave an audible warning and took itself out of its self-driving mode. But automobile emergencies take place during a fraction of a second. Google decided that while in some distant future that might be a solvable problem, it wasn't possible to solve now with existing technology.

A number of other automakers are already attempting to deal with the problem of driver distraction. Lexus and Mercedes have commercialized technology that watches the driver's eyes and head position to determine if they are drowsy or distracted. Audi in 2014 began developing a system that would use two cameras to detect when a driver was inattentive and then bring the car to an abrupt halt, if needed.

For now, however, Google seems to have changed its strategy and is trying to solve another,

perhaps simpler problem. In May of 2014, just weeks after they had given reporters an optimistic briefing on the progress of their driverless car, they shifted gears and set out to explore a new limited but more radical solution to autonomous transportation in urban environments. Unable to solve the distracted human problem, Google's engineers decided to take humans out of the loop entirely. The company de-emphasized its fleet of Prius and Lexus autonomous vehicles and set out to create a new fleet of a hundred experimental electric vehicles that entirely dispense with the standard controls in a modern automobile. Although it had successfully kept it a secret, Google had actually begun its original driverless vehicle program by experimenting with autonomous golf carts on the Google campus very early in the self-driving program. Now it was planning to return to its roots and once again autonomously ferry people around the Google campus, this time with its new specially designed vehicle fleet. Riding in the new Google car of the future will be more like riding in an elevator. The two-seat vehicle looks a bit like an ultracompact Fiat 500 or Mercedes-Benz Smart car, but with the steering wheel, gas pedal, brake, and gearshift all removed. The idea is that in crowded downtowns or on campuses passengers could enter the desired destination on their smartphones to summon a car on demand. Once they are inside, the car provides the passengers only a Trip Start button and a red E-Stop panic button. One of the conceptual shifts the engineers made was limiting the speed of the vehicle to just twenty-five miles per hour, allowing the Google cars to be regulated like golf carts rather than conventional automobiles. That meant that they could forgo air bags and other design restrictions that add cost, weight, and complexity. These limitations, however, mean that the new cars are suited only for low-speed urban driving.

Although 25 miles per hour is below highway standards, the average traffic speeds in San Francisco and New York are 18 and 17 miles per hour, respectively, and so it is possible that slow but efficient fleets of automated cars might one day replace today's taxis. A study by the Earth Institute found that Manhattan's 13,000 taxis make 470,000 trips a day. Their average speed is 10 to 11 miles per hour, carrying an average of 1.4 passengers an average distance of 2 miles, with a 5-minute average wait time to get a taxi. In comparison, the report said, it would be possible for a futuristic robot fleet of 9,000 automated vehicles hailed by smartphone to match that capacity with a wait time of less than 1 minute. Assuming a 15 percent profit, the current cost of today's taxi service would be about \$4 per trip mile, while in contrast, it was estimated, a future driverless vehicle fleet would cost about 50 cents per mile. The report showed similar savings in two other case studies—in Ann Arbor, Michigan, and Babcock Ranch, a planned community in Florida.<sup>8</sup>

Google executives and engineers have made the argument that has long been advocated by urban planners: A vast amount of space is wasted on an automotive fleet that is rarely used. Cars used to commute, for example, sit parked for much of the day, taking up urban space that could be better used for housing, offices, or parks. In urban areas, automated taxis would operate continuously, only returning to a fast-charging facility for robots to swap out their battery packs. In this light, it is easy to reimagine cities not built around private cars, with more green space and broad avenues—which would be required to safely accommodate pedestrians and cyclists.

Thrun evoked both the safety issues and the potential to redesign cities when he spoke about the danger and irrationality of our current transportation system. In addition to squandering a great deal of resources, our transportation infrastructure is responsible for more than thirty thousand road deaths annually in the United States, and about ten times more than that in both

India and China, which amounts to more than one million annual road deaths worldwide. It is a compelling argument, but it has been greeted with pushback both in terms of liability issues and more daunting ethical questions. An argument against autonomous vehicles is that the legal system is unequipped to sort out the culpability underpinning an accident that results from a design or implementation error. This issue speaks to the already incredibly complicated relationship between automotive design flaws and legal consequences. Toyota's struggle with claims of unintended acceleration, for example, cost the company more than \$1.2 billion in damages. General Motors also grappled with a design flaw involving sudden stops because of a faulty ignition switch that resulted in the recall of more vehicles than they manufactured in 2014, and may ultimately cost several billion dollars. Yet there is potentially a simple remedy for this challenge. Congress could create a liability exemption for self-driving vehicles, as it has done for childhood vaccines. Insurance companies could impose a no-fault regime when only autonomous vehicles are involved in accidents.

Another aspect of the liability issue is what has been described as a version of the "trolley problem," which is generally stated thus: A runaway trolley is hurtling down the tracks toward five people who will be killed if it proceeds on its present course. You can save these five people by diverting the trolley onto a different set of tracks that has only one person on it, but that person will be killed. Is it morally permissible to turn the trolley and thus prevent five deaths at the cost of one? First posed as a thought problem in a paper about the ethics of abortion by British philosopher Philippa Foot in 1967, it has led to endless philosophical discussions on the implications of choosing the lesser evil.<sup>9</sup> More recently it has been similarly framed for robot vehicles deciding between avoiding five schoolchildren who have run out onto the road when the only option is swerving onto the sidewalk to avoid them, thus killing a single adult bystander.

Software could generally be designed to choose the lesser evil; however, the framing of the question seems wrong on other levels. Because 90 percent of road accidents result from driver error, it is likely that a transition to autonomous vehicles will result in a dramatic drop in the overall number of injuries and deaths. So, clearly the greater good would be served even though there will still be a small number of accidents purely due to technological failures. In some respects, the automobile industry has already agreed with this logic. Air bags, for example, save more lives than are lost due to faulty air bag deployments.

Secondly, the narrow focus of the question ignores how autonomous vehicles will probably operate in the future, when it is highly likely that road workers, cops, emergency vehicles, cars, pedestrians, and cyclists will electronically signal their presence to each other, a feature that even without complete automation should dramatically increase safety. A technology known as V2X that continuously transmits the location of nearby vehicles to each other is now being tested globally. In the future, even schoolchildren will be carrying sensors to alert cars to their presence and reduce the chance of an accident.

It's puzzling, then, that the philosophers generally don't explore the trolley problem from the point of view of the greater good, but rather as an artifact of individual choice. Certainly it would be an individual tragedy if the technology fails—and of course it will fail. Systems that improve the overall safety of transportation seem vital, even if they aren't perfect. The more interesting philosophical conundrum is over the economic, social, and even cultural consequences of taking humans out of the loop in driving. More than 34,000 people died in 2013 in the United States in

automobile accidents, and 2.36 million were injured. Balance that against the 3.8 million people who earned a living by driving commercially in the United States in 2012.<sup>10</sup> Driverless cars and trucks would potentially displace many if not most of those jobs as they emerge during the next two decades.

Indeed, the question is more nuanced than one narrowly posed as a choice of saving lives or jobs. When Doug Engelbart gave what would later be billed as “The Mother of All Demos” in 1968—a demonstration of the technologies that would lead to personal computing and the Internet—he implicitly adopted the metaphor of driving. He sat at a keyboard and a display and showed how graphical interactive computing could be used to control computing and “drive” through what would become known as cyberspace. The human was very much in control in this model of intelligence augmentation. Driving was the original metaphor for interactive computing, but today Google’s vision has changed the metaphor. The new analogy will be closer to traveling in an elevator or a train without human intervention. In Google’s world you will press a button and be taken to your destination. This conception of transportation undermines several notions that are deeply ingrained in American culture. In the last century the car became synonymous with the American ideal of freedom and independence. That era is now ending. What will replace it?

It is significant that Google is instrumental in changing the metaphor. In one sense the company began as the quintessential intelligence augmentation, or IA, company. The PageRank algorithm Larry Page developed to improve Internet search results essentially mined human intelligence by using the crowd-sourced accumulation of human decisions about valuable information sources. Google initially began by collecting and organizing human knowledge and then making it available to humans as part of a glorified Memex, the original global information retrieval system first proposed by Vannevar Bush in the *Atlantic Monthly* in 1945.<sup>11</sup>

As the company has evolved, however, it has started to push heavily toward systems that replace rather than extend humans. Google’s executives have obviously thought to some degree about the societal consequences of the systems they are creating. Their corporate motto remains “Don’t be evil.” Of course, that is nebulous enough to be construed to mean almost anything. Yet it does suggest that as a company Google is concerned with more than simply maximizing shareholder value. For example, Peter Norvig, a veteran AI scientist who has been director of research at Google since 2001, points to partnerships between human and computer as the way out of the conundrum presented by the emergence of increasingly intelligent machines. A partnership between human chess experts and a chess-playing computer program can outplay even the best AI chess program, he notes. “As a society that’s what we’re going to have to do. Computers are going to be more flexible and they’re going to do more, and the people who are going to thrive are probably the ones who work in a partnership with machines,” he told a NASA conference in 2014.<sup>12</sup>

What will the partnerships between humans and intelligent cars of the future look like? What began as a military plan to automate battlefield logistics, lowering costs and keeping soldiers out of harm’s way, is now on the verge of reframing modern transportation. The world is plunging ahead and automating transportation systems, but the consequences are only dimly understood today. There will be huge positive consequences in safety, efficiency, and environmental quality. But what about the millions of people now employed driving throughout the world? What will they do when they become the twenty-first-century equivalent of the blacksmith or buggy-whip

maker?